

A Metadata Calculus for Securing Information Flows

Mudhakar Srivatsa and Dakshi Agrawal
IBM T.J. Watson Research Center
{msrivats, agrawal}@us.ibm.com

Shane Balfe
Royal Holloway, University of London
s.balfe@rhul.ac.uk

Abstract

Traditional approaches to information sharing use a highly conservative approach to deduce the metadata for an output object x derived from input objects y_1, y_2, \dots, y_n (e.g.: maximum over the security labels of all input objects). Such approaches does not account for functions that explicitly downgrade the value of an object. Consequently, the security labels in traditional approaches tend to monotonically increase as newer objects are derived from existing ones. In this paper we present a novel metadata calculus for securing information flows. The metadata calculus defines a metadata vector space that supports a time varying value function that is computed as a function of the object's metadata and operators $+$ and \cdot to compute the metadata of an output object that is derived by downgrading, transforming or fusing other objects. We also describe a concrete realization of our metadata calculus wherein the tightness of our value estimates competes in an optimization problem. We present several tradeoffs with space and accuracy and explore a spectrum of solutions ranging from conservative to risk-based value estimates.

1 Introduction

Large corporations are slowly being transformed from monolithic, vertically integrated entities, into globally disaggregated value networks, where each member focuses on its core competencies and relies on partners and suppliers to develop and deliver goods and services. The ability of multiple partners to come together, share sensitive business information and coordinate activities to rapidly respond to business opportunities, is becoming a key driver for success.

The defense sector too, has similar, dynamic information sharing needs. Traditional wars between

armies of nation-states are being replaced by highly dynamic missions where teams of soldiers, strategists, logisticians, and support staff, drawn from a coalition of military organizations as well as local (military and civilian) authorities, fight against elusive enemies that easily blend into the civilian population [6]. Securely disseminating mission critical tactical intelligence to the pertinent people in a timely manner will be a critical factor in a mission's success.

While it is clear that information sharing across organizational boundaries is becoming a necessity, it is important for the recipient to ensure that it receives high quality information from the sender. However, for a sender to share high quality information, the sender needs assurance from the recipient that the shared information will not be *misused* (e.g.: unregulated or unintended information disclosure). Poor quality of information and unauthorized information disclosure can create the risk of legal liability, financial loss, tarnished reputation, or in some environments, a loss of life. Evidently, there is a risk related tradeoff between the quality of information and information misuse. Understanding this tradeoff *minimally* requires us to quantify the *value* of information being shared.

Unfortunately, traditional approaches to information sharing suffer from two major drawbacks. First, they use fairly static security labels to tag information, and thus do not attempt to capture dynamic attributes of tactical information such as time sensitivity, accuracy, etc. The value of a piece of information (henceforth, called an object) is computed as a function of its security labels (henceforth, called metadata). For example, Multi-Level Security (MLS) labels such as unclassified, classified, secret, top secret are used to enforce mandatory access control in a military setting; Decentralized Label Management (DLM) labels each object with allow and deny lists and regulates information flows

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE DEC 2008		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE A Metadata Calculus for Securing Information Flows				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) IBM T.J. Watson Research Center				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM002187. Proceedings of the Army Science Conference (26th) Held in Orlando, Florida on 1-4 December 2008					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 7	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

	MLS	DLM
Metadata of x	l_x	A_x, D_x
Metadata of $g(x_1, \dots, x_n)$	$\max(l_{x_1}, \dots, l_{x_n})$	$\bigcap_{i=1}^n A_{x_i}, \bigcup_{i=1}^n D_{x_i}$

Figure 1: Metadata Calculus

based on these labels.

Second, traditional approaches to information sharing use a highly conservative approach to deduce the metadata for an output object x derived from input objects x_1, x_2, \dots, x_n (see Figure 1). Such approaches does not account for functions that explicitly downgrade the value of an object. For example, a set of numeric objects may be statistically downgraded using their mean; an image may be downgraded by lowering its resolution, smoothness factor, etc. Consequently, the security labels in traditional approaches tend to monotonically increase as newer objects are derived from existing ones. Ultimately, most of the derived objects receive the highest security label making them inaccessible to most of the users in the organization. As pointed out by the JASON committee report [5], traditional approaches severely constrict the flow of critical information and are thus not appropriate for dynamic settings, where systems and processes evolve rapidly and there are transient needs for sharing tactical, time-sensitive information across organizational boundaries.

In this paper we present a novel metadata calculus for securing information flows. We capture an object’s metadata using a vector space \mathcal{M} and define a metadata calculus that supports the following primitives: (i) a time varying value function Γ based on an object’s metadata, and (ii) operators $+$ and \cdot on the vector space \mathcal{M} that can be used to compute the metadata for an output object that is derived by downgrading, transforming or fusing other objects. We also describe a concrete realization of our metadata calculus wherein the tightness of our value estimates competes in an optimization problem. We present several tradeoffs with space and accuracy and explore a spectrum of solutions ranging from conservative to risk-based value estimates.

The rest of this paper is organized as follows. Section 2 describes related work on risk based secure information flows. Section 3 describes metadata types and presents an information theoretic

approach to estimating the value of an object. Section 4 describes a calculus for succinctly computing the metadata for an output object that is derived from one or more input objects. Finally we conclude the paper in Section 5.

2 Related Work

There has been significant research on decentralized information labels and assured information sharing within and across multiple organizations [2, 3, 4, 10, 11, 12] in recent years. However, these works primarily focus on the problem of specifying and manipulating the sharing, propagation and downgrading constraints on data. These works also assume appropriate security controls that manipulate, bind and respect these labels are already in place, for example, via a secure distributed runtime language, or some other form of a secure distributed trusted computing base. Clearly, in practice, for the settings described in the introduction, one partner cannot be sure of either the existence, or the proper usage of, a secure runtime environment of another partner.

Recently, new approaches based on risk estimation and economic mechanisms have been proposed for enabling the sharing of information in dynamic environments [1, 5]. These approaches are based on the idea that the sender dynamically computes an estimate of the risk of information disclosure in providing information to a receiver based on the secrecy of the information to be divulged and the sender’s estimate on the trustworthiness of the recipient. The sender then “charges” the receiver for this estimated risk. The recipient, in turn, can decide which type of information is most useful to him and pay only to access that. Entities would either be given a line of risk credit, or adopt a market-based mechanism to “purchase risk” using a pseudo-currency. Under the assumption that the line of risk credit or the risk available for purchase in the market is limited, an entity will be encouraged to be frugal with their amassed risk credits and, con-

sequently, reluctant to spend them unnecessarily. Since all information flows are “charged” against expected losses due to unauthorized disclosure and the amount of risk available is limited, an argument is made that the total information disclosure risk incurred by an organization is controlled.

While, as a concept, using risk estimation, charging for risk of information flows, and limited risk credits are promising ideas for enabling information sharing in dynamic environments, the existing work in this area [1, 5] has gaps in how this concept can be realized to enable cross-organizational secure information flows in dynamic environments such as between organizations or partners in a coalition. In [1, 5, 8, 7], while risk is estimated based on the object metadata [9], the actual formulas or examples use static credentials (e.g., the security clearance or category set) of the recipient, rather than a dynamic value of the object. Indeed, the value of most tactical information tends to decrease with time and evolve as the object is downgraded, transformed or fused with other objects. In this paper we present a novel metadata calculus that can be used to succinctly estimate the time varying value of tactical information in a dynamic coalition setting.

3 Metadata Model

In this section we describe our metadata model. For the sake of simplicity we include one dynamic attribute, namely time, in our metadata model. We represent the metadata for an object x as a vector \vec{x} in a vector space \mathcal{M} . \mathcal{M} also supports an unary operator that maps $\vec{x} \in \mathcal{M}$ to value function: $\Gamma: \mathcal{M} \rightarrow (\mathcal{F} \rightarrow \mathcal{F})$, where \mathcal{F} denotes an integer or real number field. For example, $\vec{x} = (10, 2)$ and $\Gamma\vec{x}(t) = \max(10 - 2t, 0)$ or $\Gamma\vec{x}(t) = 10e^{-2t}$. The value operator Γ satisfies the following properties:

$$\begin{aligned} 0 &\leq \Gamma\vec{x}(t) < \infty, \forall t \\ \frac{\partial \Gamma\vec{x}(t)}{\partial t} &\leq 0, \forall t \\ x \subseteq y &\Rightarrow \Gamma\vec{x}(t) \leq \Gamma\vec{y}(t), \forall t \end{aligned}$$

Constraining the value of object to non-negative integers (or real numbers) may be questionable. One can think of sources of disinformation (misguiding information) to have a negative value. In this paper we do not consider pieces of information that are intended to misguide the recipient. In the absence of

disinformation, the value of information is monotonic, that is, if an object x is completely contained in object y , then $\Gamma\vec{x}(t) \leq \Gamma\vec{y}(t)$.

The value of an output object x computed as $g(y_1, y_2, \dots, y_n)$ is computed as shown by an empirical formula in Equation 1, where $\bar{Y}_i = \{Y_1, \dots, Y_{i-1}, Y_{i+1}, \dots, Y_n\}$ and $\bar{y}_i = \{y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_n\}$.

$$\Gamma\vec{x}(t) = \sum_{i=1}^n \Gamma\vec{y}_i(t) * \frac{f_{Y_i|X}(y_i|x, B)}{f_{X|\bar{Y}_i}(x|\bar{y}_i)} \quad (1)$$

We use f_X to denote the probability distribution function for a random variable X . Value computation uses the notion of self-information expressed as $I(y_i|x) = D(\delta_{y_i} \parallel f_{Y_i|X}(y_i|x)) = -\log(f_{Y_i|X}(y_i|x))$, where $D(X \parallel Y)$ denotes KL-divergence between probability distributions X and Y and δ_{y_i} denotes the Dirac delta function whose value is one when $Y_i = y_i$ and zero otherwise. Intuitively self-information $I(y_i|x)$ denotes the number of additional bits that need to be learnt in order to reconstruct y_i given that the entity knows the probability distribution $f_{Y_i|X}$. Hence, $2^{-I(y_i|x)} = f_{Y_i|X}(y_i|x)$ denotes the fraction of information about y_i that may be inferred from x . We remark that exact reconstruction of y_i may not be required for certain objects (e.g.: geographical location). In such cases, one can replace δ_{y_i} by some probability distribution that is centered around y_i .

We argue that Equation 1 satisfies the intuitive notion of object downgrade, transforms and fusion. In the rest of this section, we demonstrate the applicability of Equation 1 to a wide range of functions $g(\cdot)$ ranging from arithmetic functions, database operations and cryptographic functions. Figure 2, 3 and 4 show value computations for some sample functions g . We use B to denote background information known to the consumer of object x such as cryptographic secrets.

In the case of bijective arithmetic functions (such as $x = g(y_1) = y_1 + 1$), we note that given x and the function g , one can completely recover all information about y_1 . Hence, the value of x equals the value of y_1 for all time instances t . On the other hand, arithmetic functions such as $x = y_1^2$ loose information on y_1 ; in particular, given x one can identify two possible values for y_1 (namely, $\pm\sqrt{x}$). In the absence of any background information on y_1 , this results in an entropy loss of one bit; equivalently

g	B	$\Gamma \vec{x}$	note
$x = y_1 + 1$	—	$\Gamma \vec{y}_1(t)$	g is bijective
$x = y_1^2$	—	$\frac{\Gamma \vec{y}_1(t)}{2}$	g is invertible
$x = y_1^2$	$y_1 > 0$	$\Gamma \vec{y}_1(t)$	g is bijective given B
$x = \sum_{i=1}^n y_i$	$n, \forall i, y_i \sim f_Y$	$0 \leq \Gamma \vec{x}(t) \leq \sum_{i=1}^n \Gamma \vec{y}_i(t)$	Statistical downgrading

Figure 2: Value Computation for Arithmetic Functions

g	B	$\Gamma \vec{x}$	note
$x = y_1 \bowtie y_2$	—	$2^k * (\Gamma \vec{y}_1(t) + \Gamma \vec{y}_2(t))$	Join operation
$x = y_1 \cup y_2$	—	$\Gamma \vec{y}_1(t) + \Gamma \vec{y}_2(t)$	Union operation
$x = \pi_{y_1} \{y_1, y_2\}$	—	$2^{-k} * \overrightarrow{\Gamma \{y_1, y_2\}}$	Project operation
$x = \sigma_p \{y_1, y_2\}$	—	$\frac{ x }{ \{y_1, y_2\} } * \overrightarrow{\Gamma \{y_1, y_2\}}$	Select operation using predicate p

Figure 3: Value Computation for Database Functions: Assume, we have two database fields Y_1 and Y_2 . The values y_1 in Y_1 and y_2 in Y_2 are independent and chosen using a uniform distribution between $(0, 2^k - 1)$

g	B	$\Gamma \vec{x}$	note
$x = K$	—	0	Key K has zero value; but knowledge of K affects the value of other objects
$x = E_K(y_1)$	—	0	Ideal encryption
$x = E_K(y_1)$	K	$\Gamma \vec{y}_1(t)$	Ideal encryption
$x = D_{K'}(E_K(y_1))$	$K' \neq K$	0	Ideal decryption
$x = D_K(E_K(y_1))$	K	$\Gamma \vec{y}_1(t)$	Ideal decryption
$x = H_K(y_1)$	K	0	Ideal hash
$x = Sig_K(y_1)$	K	0	Ideal signature

Figure 4: Value Computation for Cryptographic Functions

the value of x is half the value of y_1 . However, the recipient were to know that $y_1 \geq 0$, then there is no loss of information. Functions such as sum and average may exhibit different information loss characteristics. For example, let us suppose $x = y_1 + y_2$. If the recipient knows that $0 \leq y_1, y_2 \leq 5$, then given $x = 0$ (or 10), it can obtain all information about y_1 and y_2 respectively.

For database operations, we recognize the need to differentiate between union and join operations. Union operates on two sets of the *same type*; for example, let us consider sets of type color $y_1 = \{\text{red}\}$ and $y_2 = \{\text{blue}\}$. Consequently, union of two sets does not result in any additional information than the input sets. On the other hand, join operates on two sets of *different types*; for example, let us consider a set of type $\langle x\text{-coord}, \text{id} \rangle$ $y_1 = \{\langle 10, \text{id}_1 \rangle\}$ and a set of type $\langle y\text{-coord}, \text{id} \rangle$ $y_2 = \{\langle 5, \text{id}_1 \rangle\}$. A join on y_1 and y_2 reveals the (x, y) coordinates of the entity (id_1) and thus has significantly more information than the input sets. To cite another example, let us consider a 128-bit cryptographic key $K = L \parallel R$ of value V , where L and R denote the left and the right 64-bites of a 128-bit key K . One might argue that the value of L and R is $2^{-64} * V$, assuming the key K is randomly chosen over a 128-bit field. Now, let us consider a join of sets of type $\langle L, \text{kid} \rangle$ and $\langle R, \text{kid} \rangle$, where kid denotes key identifier: $y_1 = \{\langle L_1, \text{kid}_1 \rangle\}$ and $y_2 = \{\langle R_1, \text{kid}_1 \rangle\}$. It is easy to see that the value of x must be significantly higher than the sum of values y_1 and y_2 . Indeed, Equation 1 amplifies the values of y_1 and y_2 by a factor of 2^{64} when deriving the value of x .

For cryptographic operations, we model ideal behavior using a 0/1 value relationship with the input object. For example, using an ideal encryption function $x = E_K(y_1)$, the value of x is zero if the recipient does not know K ; otherwise, the value of x is equal to the value of y_1 , since the recipient can recover all information about y_1 using the corresponding decryption function D and the key K .

4 Metadata Calculus

In this section, we describe a metadata calculus using two binary operators on the metadata vector space \mathcal{M} : vector addition $+$: $\mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}$ and scalar multiplication \cdot : $\mathcal{F} \times \mathcal{M} \rightarrow \mathcal{M}$, where \mathcal{F} denotes a field such as integers or real numbers. These




Figure 5: Conservative Vs Risky ‘+’ Operator

binary operators satisfy the following homomorphic properties:

- $\vec{x} = \vec{y}_1 + \vec{y}_2 \Leftrightarrow \Gamma \vec{x}(t) = \Gamma \vec{y}_1(t) + \Gamma \vec{y}_2(t)$ for all t .
- $\vec{x} = a \cdot \vec{y}_1 \Leftrightarrow \Gamma \vec{x}(t) = a * \Gamma \vec{y}_1(t)$ for all t .

In addition, they also satisfy the following intuitive properties:

- *Commutative*: For any $\vec{y}_1, \vec{y}_2 \in \mathcal{M}$, $\vec{y}_1 + \vec{y}_2 = \vec{y}_2 + \vec{y}_1$.
- *Associative*: For any $\vec{y}_1, \vec{y}_2, \vec{y}_3 \in \mathcal{M}$, $\vec{y}_1 + (\vec{y}_2 + \vec{y}_3) = (\vec{y}_1 + \vec{y}_2) + \vec{y}_3$.
- *Zero vector* $\vec{0}$: For any $\vec{y}_1 \in \mathcal{M}$, $\vec{y}_1 + \vec{0} = \vec{y}_1$.
- *Distributive over + in \mathcal{M}* : For any $a \in \mathcal{F}$, $\vec{y}_1, \vec{y}_2 \in \mathcal{M}$, $a \cdot (\vec{y}_1 + \vec{y}_2) = a \cdot \vec{y}_1 + a \cdot \vec{y}_2$.
- *Distributive over + in \mathcal{F}* : For any $a, b \in \mathcal{F}$, $\vec{y}_1 \in \mathcal{M}$, $(a + b) \cdot \vec{y}_1 = a \cdot \vec{y}_1 + b \cdot \vec{y}_1$.
- *Distributive over \cdot in \mathcal{M}* : For any $a, b \in \mathcal{F}$, $\vec{y}_1 \in \mathcal{M}$, $a \cdot (b \cdot \vec{y}_1) = (ab) \cdot \vec{y}_1$.
- *Scalar 1 in \mathcal{F}* : For any $\vec{y}_1 \in \mathcal{M}$, $1 \cdot \vec{y}_1 = \vec{y}_1$.

Based on the properties described above, it is easy to see that when x is computed as $g(y_1, y_2, \dots, y_n)$ then the metadata \vec{x} can be computed as shown in Equation 2. Indeed give the homomorphic properties on the \mathcal{M} , we can show that Equation 2 implies Equation 1 for all time t .

$$\vec{x} = \sum_{i=1}^n \frac{f_{Y_i|X}(y_i|x, B)}{f_{X|\vec{Y}_i}(x|\vec{y}_i)} \cdot \vec{y}_i \quad (2)$$

4.1 Realizing the Metadata Calculus

In this section, we describe a concrete instantiation of metadata calculus. Unfortunately, there exists no vector space that satisfies all the required homomorphic properties. In this section, we describe a metadata vector space and value functions that satisfy a weaker notion of homomorphism as shown in Equation 3. It follows from Equation 3 that our metadata calculus makes conservative estimates on object values.

$$\begin{aligned} \vec{x} &= a \cdot \vec{y}_1 \Leftrightarrow \Gamma \vec{x}(t) = a * \Gamma \vec{y}_1(t), \forall t \\ \vec{x} &= \vec{y}_1 + \vec{y}_2 \Rightarrow \Gamma \vec{x}(t) \geq \Gamma \vec{y}_1(t) + \Gamma \vec{y}_2(t), \forall t \end{aligned} \quad (3)$$

We use a metadata space $\mathcal{M} = \mathcal{Z} \times \mathcal{Z}$, where \mathcal{Z} denotes integer field. Given $\vec{x} = (c_0, c_1)$, $\Gamma \vec{x}(t) = \max(c_0 - c_1 * t, 0)$. Given $\vec{x}_1 = (c_0^1, c_1^1)$ and $\vec{x}_2 = (c_0^2, c_1^2)$, the $+$ and \cdot operators are defined as follows:

$$\begin{aligned} a \cdot \vec{x}_1 &= (a * c_0^1, a * c_1^1) \\ \vec{x}_1 + \vec{x}_2 &= \left(c_0^1 + c_0^2, \frac{c_0^1 + c_0^2}{\max\left(\frac{c_0^1}{c_1^1}, \frac{c_0^2}{c_1^2}\right)} \right) \end{aligned}$$

Figure 5 illustrates the $+$ operator on $\vec{y}_1 = (10, 2)$ and $\vec{y}_2 = (20, 10)$. It is easy to see that $\Gamma \vec{y}_1(t) + \Gamma \vec{y}_2(t)$ is:

$$\Gamma \vec{y}_1(t) + \Gamma \vec{y}_2(t) = \begin{cases} 30 - 12t & \text{if } t \leq 2 \\ 10 - 2t & \text{if } 2 < t \leq 5 \end{cases}$$

It is easy to see that the above equation cannot be represented by a straight line and thus cannot be mapped into a metadata vector in \mathcal{M} . Hence, we choose the least conservative straight line $(30 - 6t)$ such that $\Gamma \vec{y}_1(t) + \Gamma \vec{y}_2(t) \leq 30 - 6t$ for all t . Indeed, there are several options to ensure that our value estimates are tighter all of which can be modeled as optimization problems.

- First, we can increase the dimensionality of the metadata vector space and use a high order polynomial for $\Gamma \vec{x}(t)$. It is easy to see that proposed metadata calculus can be extended to all value functions that are polynomial in time t . We note that increasing the dimensionality of the metadata vector allows us to compute tighter value estimates at the cost of higher storage cost.

- Second, we can represent \vec{x} as k tuples where each tuple describes a straight line within some time interval. We note that when x is computed as a function of y_1, \dots, y_n , then the value of x may be represented by at most n linear constraints. For any given constant $k \leq n$, we can compute a set of k linear constraints that tightly bounds the set of n linear constraints. We note that increasing k allows us to compute tighter value estimates at the cost of higher storage cost.
- Third, we could permit bounded violations to the constraint $\Gamma(\vec{y}_1 + \vec{y}_2)(t) \geq \Gamma \vec{y}_1(t) + \Gamma \vec{y}_2(t)$ for some instants t . One can quantify the risk (r^+) in a value estimate using the area enclosed by the region wherein our value estimate is lower than the true value of the object; similarly, one can quantify overestimation (r^-) using the area enclosed by the region wherein our value estimate is higher than the true value of the object. We formulate two optimization problems that allow us to trade off the conservativeness and tightness in our estimates. First, we can restrict the risk in our value estimate to at most $X\%$ of the value of the object (averaged over its lifetime). Second, we can attempt to minimize a function of risk and overestimation, say $\alpha * r^+ - r^-$, for some $\alpha > 0$.

5 Conclusion

In this paper we have presented a novel metadata calculus for securing information flows in a tactical setting. The metadata calculus defines a metadata vector space that supports a time varying value function that is computed as a function of the object's metadata and operators $+$ and \cdot to compute the metadata of an output object that is derived by downgrading, transforming or fusing other objects. We have also described a concrete realization of our metadata calculus using a value function that is polynomial in time t . We have formulated the problem of finding tight value estimates as various optimization problems. These formulations model various tradeoffs with space and accuracy and explore a spectrum of solutions ranging from conservative to risk-based value estimates.

Acknowledgements

Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorised to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

- [1] P.-C. Cheng, P. Rohatgi, C. Keser, P.A. Karger, G.M. Wagner, and A.S. Reninger. Fuzzy Multi-Level Security: An Experiment on Quantified Risk-Adaptive Access Control. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy (SP 2007)*, pages 222–230. IEEE Computer Society, 2007.
- [2] C.D. McCollum and J.R. Messing L. Notargiacomo. Beyond the Pale of MAC and DAC-Defining New Forms of Access Control. In *Proceedings of the 1990 IEEE Symposium on Security and Privacy (S&P 1990)*, pages 190–200. IEEE Computer Society, 1990.
- [3] A.C. Myers and B. Liskov. A Decentralized Model for Information Flow Control. In *Proceedings of the 1997 Symposium on Operating Systems Principles (SOSP 1997)*, pages 129–142. ACM Press, 1997.
- [4] A.C. Myers and B. Liskov. Complete Safe Information Flow with Decentralized Labels. In *Proceedings of the 1998 IEEE Symposium on Security and Privacy (S&P 1998)*, pages 186–197. IEEE Computer Society, 2001.
- [5] Jason Program Office. HORIZONTAL INTEGRATION: Broader Access Models for Realizing Information Dominance. Special Report JSR-04-13, MITRE Corporation, 2004.
- [6] D. Roberts, G. Lock, and D.C. Verma. Holistan: A Futuristic Scenario for International Coalition Operations. In *In Proceedings of Fourth International Conference on Knowledge Systems for Coalition Operations (KSCO 2007)*, 2007.
- [7] M. Srivatsa, P. Rohatgi, and S. Balfe. Securing information flows: A quantitative risk analysis approach. In *Proceedings of 15th ACM Conference on Computer and Communication Security (CCS)*, 2008.
- [8] M. Srivatsa, P. Rohatgi, S. Balfe, and K. G. Paterson. Trust management for secure information flows. In *Proceedings of 27th IEEE Conference on Military Communications (MILCOM)*, 2008.
- [9] M. Srivatsa, P. Rohatgi, S. Balfe, and S. Reidt. Securing information flows: A metadata framework. In *Proceedings of 1st IEEE Workshop on Quality of Information for Sensor Networks (QoISN)*, 2008.
- [10] N. Swamy, M. Hicks, and S. Tsang. Verified Enforcement of Security Policies for Cross-Domain Information Flows. In *Proceedings of the 2007 Military Communications Conference (MILCOM 2007)*, pages 192–206. IEEE Computer Society, 2007.
- [11] J.A. Vaughan and S. Zdancewic. A Cryptographic Decentralized Label Model. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy (S&P 2007)*, pages 192–206. IEEE Computer Society, 2007.
- [12] S. Zdancewic and A.C. Myers. Secure Information flows and CPS. In *Proceedings of the 10th European Symposium on Programming (ESOP 2001)*, pages 46–61. Springer, 2001.